

1 access control entries in a manner analogous to the read access control discussed
2 herein.

3 With convergent encryption, one encrypted version of the file is stored and
4 replicated among the serverless distributed file system 150. Along with the
5 encrypted version of the file is stored one or more access control entries depending
6 upon the number of authorized users who have access. Thus, a file in the
7 distributed file system 150 has the following structure:

$$[E_{h(F)}(F), \langle E_{K1}(h(F)) \rangle, \langle E_{K2}(h(F)) \rangle, \dots, \langle E_{Km}(h(F)) \rangle]$$

11 One advantage of convergent encryption is that the encrypted file can be
12 evaluated by the file system to determine whether it is identical to another file
13 without resorting to any decryption (and hence, without knowledge of any
14 encryption keys). Unwanted duplicative files can be removed by adding the
15 authorized user(s) access control entries to the remaining file. Another advantage
16 is that the access control entries are very small in size, on the order of bytes as
17 compared to possibly gigabytes for the encrypted file. As a result, the amount of
18 overhead information that is stored in each file is reduced. This enables the
19 property that the total space used to store the file is proportional to the space that is
20 required to store a single encrypted file, plus a constant amount of storage for each
21 additional authorized reader of the file.

22 For more information on convergent encryption, the reader is directed to
23 co-pending U.S. Patent Application Serial No. 09/565,821, ^{still pending} entitled "Encryption
24 Systems and Methods for Identifying and Coalescing Identical Objects Encrypted
25 with Different Keys", which was filed May 5, 2000, in the names of Douceur et

211
6/2/05

1 al., and is commonly assigned to Microsoft Corporation. This application is
2 hereby incorporated by reference.

3 For small files, the entire file is hashed and encrypted using convergent
4 encryption, and the resulting hash value is used as the encryption key. The
5 encrypted file can be verified without knowledge of the key or any need to decrypt
6 the file first. For large files, the file contents are broken into smaller blocks and
7 then convergent encryption is applied separately to each block. For example, the
8 file F may be segmented into "n" pages F^0 - F^{n-1} , where each page is a fixed size
9 (e.g., a 4Kbyte size). Convergent encryption is then applied to the file at the block
10 level. That is, each block F^i is separately hashed using a one-way hash function
11 (e.g., SHA, MD5, etc.) to produce a hash value $h(F^i)$. Each block F^i is then
12 encrypted using a symmetric cipher (e.g., RC4, RC2, etc.) with the hash value
13 $h(F^i)$ as the key, or $E_{h(F^i)}(F^i)$, resulting in an array of encrypted blocks which form
14 the contents of the file. For more information on block-by-block encryption, the
15 reader is directed to co-pending U.S. Patent Application Serial No. 09/814,259 entitled
16 "On-Disk File Format for Serverless Distributed File System", ~~Attorney Docket~~
17 ~~No. MS1-733US~~, to inventors William J. Bolosky, Gerald Cermak, Atul Adya, and
18 John R. Douceur. This application is hereby incorporated by reference.

19 File information generation module 220 can generate the file information at
20 any of a wide variety of times. In one implementation, module 220 is designed to
21 operate as a background process. When files are created or modified, the file
22 names are added to a queue to be acted on by module 220. When computing
23 device 200 is not busy (e.g., the processor has free cycles, or has been idle for a
24 period of time), module 220 operates to generate file information for one of the
25 files in the queue. Alternatively, module 220 may be designed to run at times of

pending

JH
6/15/05